CLAIMS

1.   A method for selecting master nodes to manage a target node group in a computer network having multiple nodes and node groups, the method comprising:

determining a hamming distance for node pairs within the network, each of the node pairs having two node pair members and the hamming distance being the number of node groups the node pair members do not share in common;

determining a participation index for nodes within the network, the participation index being the number of node groups a node belongs to;

determining an availability potential for the node pairs, the availability potential being the sum of the participation indexes of the node pair members subtracted by the hamming distance of the node pair;

finding an optimal combination of node pairs, the optimal combination of node pairs having a maximum total availability potential for the network;

selecting a master node pair for the target node group, the master node pair being the node pair from the optimal combination of node pairs having both node pair members belonging to the target node group; and

if the master node pair does not exist for the target node group, selecting a master node for the target node group, the master node being the node belonging to the target node group.

2.   The method of claim 1, wherein either of the node pair members of the master node pair selected for the target node group are not master nodes for another node group with a third node.

3.   The method of claim 1, wherein selecting a master node pair for the target node group further includes if the optimal combination of node pairs contains more than one node pair having both node pair members belonging to the target node group, selecting the node pair

IBM Docket No. SJO920010172

5   with the smallest sum of the participation indexes of the node pair

members.

4.   The method of claim 1, wherein selecting a master node for

the target node group further includes if more than one node belongs to

the target node group, selecting the node with the smallest

participation index.

5.   The method of claim 1, wherein determining a hamming

distance for node pairs within the network includes XORing the node

pair members.

6.   In a computer-readable medium having stored thereon a data

structure for use in selecting master nodes to manage a target node

group in a computer network having multiple nodes and overlapping node

groups, the data structure comprising:

5   a hamming distance array containing, for node pairs in the

network, the number of node groups node pair members do not share in

common;

a participation index array containing, for nodes in the network,

the number of node groups the node belongs to; and

10   an availability potential array containing, for node pairs in the

network, the sum of the participation indexes of the node pair members

subtracted by the hamming distance of the node pair.

7.   The data structure of claim 6, wherein the hamming distance

array is an n by n array, where n is the number of nodes in the

network.

8.   The data structure of claim 6, wherein the participation

index array is a *one* by n array, where n is the number of nodes in the

network.

9.   The data structure of claim 6, wherein the availability

potential array is an n by n array, where n is the number of nodes in

the network.

10.   The data structure of claim 6, further comprising a node

group configuration array storing the relationships between each node
and each node group in the network.

11.   The data structure of claim 10, wherein the node group
configuration array is an m by n array, where m is the number of node
groups in the network and n is the number of nodes in the network.

12.   The data structure of claim 10, wherein in the node group
configuration array a "1" is used to signify the node belonging to the
node group and a "0" is used to signify the node not belonging to the
node group.

13.   The data structure of claim 6, further comprising an
optimal combination of node pairs, the optimal combination of node
pairs being the combination of node pairs having a maximum total
availability potential for the network.

14.   The data structure of claim 6, further comprising a
preferred order array listing a hierarchy of master node pairs for node
groups in the network.

15.   The data structure of claim 14, wherein the preferred order
array further includes a hierarchy of master nodes for node groups in
the network.

16.   A computer program product embodied in a tangible media
comprising:

computer readable program codes coupled to the tangible media for
selecting master nodes to manage a target node group in a computer
5   network having multiple nodes and node groups, the computer readable
program codes comprising:

first computer readable program code configured to cause the
program to determine a hamming distance for node pairs within the
network, the node pair having two node pair members and the hamming
10   distance being the number of node groups the node pair members do not
share in common;

second computer readable program code configured to cause the

program to determine a participation index for nodes within the network, the participation index being the number of node groups a node

15   belongs to;

third computer readable program code configured to cause the program to determine an availability potential for each node pair, the availability potential being the sum of the participation indexes of the node pair members subtracted by the hamming distance of the node

20   pair;

fourth computer readable program code configured to cause the program to find an optimal combination of node pairs, the optimal combination of node pairs having the maximum total availability potential for the network; and

25   fifth computer readable program code configured to cause the program to select a master node pair for the target node group, the master node pair being the node pair from the optimal combination of node pairs having both node pair members belonging to the target node group; and

30   if the master node pair does not exist for the target node group, sixth computer readable program code configured to cause the program to select a master node for the target node group, the master node being the node belonging to the target node group.

17.   The computer program product of claim 16, wherein the tangible media comprises a magnetic disk.

18.   The computer program product of claim 16, wherein the tangible media comprises an optical disk.

19.   The computer program product of claim 16, wherein the tangible media comprises a propagating signal.

20.   The computer program product of claim 16, wherein the tangible media comprises a random access memory device.

21.   A system for selecting master nodes to manage a target node group in a computer network having multiple nodes and node groups, the system comprising:

IBM Docket No. SJO920010172

a hamming distance module for determining a hamming distance of

5   node pairs within the network, each of the node pairs having two node

pair members and the hamming distance being the number of node groups

the node pair members do not share in common;

a participation index module for determining a participation

index of nodes within the network, the participation index being the

10   number of node groups a node belongs to;

an availability potential module for determining an availability

potential of the node pairs, the availability potential being the sum

of the participation indexes of the node pair members subtracted by the

hamming distance of the node pair;

15   a search module for finding an optimal combination of node pairs,

the optimal combination of node pairs having the maximum total

availability potential for the network;

a first selection module for selecting a master node pair of the

target node group, the master node pair being the node pair from the

20   optimal combination of node pairs having both node pair members

belonging to the target node group; and

a second selection module for selecting a master node for the

target node group if the master node pair does not exist for the target

node group, the master node being the node belonging to the target node

25   group.

22.   The system of claim 21, wherein the node pair members of

the master node pair selected for the target node group are not master

nodes for another node group with a third node.

23.   The system of claim 21, wherein the first selection module

further includes if the optimal combination of node pairs contains more

than one node pair having both node pair members belonging to the

target node group, selecting the node pair with the smallest sum of the

5   participation indexes of the node pair members.

24.   The system of claim 21, wherein the second selection module further includes if more than one node belongs to the target node group, selecting the node with the smallest participation index.

25.   The system of claim 21, wherein the hamming distance module includes XORing the node pair members.

26.   A method for selecting master nodes to manage a target node group in a computer network having multiple nodes and node groups, the method comprising:

    determining a hamming distance for node pairs within the network,
5   each of the node pairs having two node pair members and the hamming distance being the number of node groups the node pair members do not share in common;

    finding an optimal combination of node pairs, the optimal combination of node pairs having a minimum total hamming distance for
10 the network;

    selecting a master node pair for the target node group, the master node pair being the node pair from the optimal combination of node pairs having both node pair members belonging to the target node group; and

15 if the master node pair does not exist for the target node group, selecting a master node for the target node group, the master node being the node belonging to the target node group.

27.   The method of claim 26, wherein either of the node pair members of the master node pair selected for the target node group are not master nodes for another node group with a third node.

28.   The method of claim 26, wherein selecting a master node pair for the target node group further includes:

    determining a participation index for nodes within the network, the participation index being the number of node groups a node belongs
5   to; and

    if the optimal combination of node pairs contains more than one node pair having both node pair members belonging to the target node

group, selecting the node pair with the smallest sum of the participation indexes of the node pair members.

29.   The method of claim 26, wherein selecting a master node for the target node group further includes:

determining a participation index for nodes within the network, the participation index being the number of node groups a node belongs to; and

if more than one node belongs to the target node group, selecting the node with the smallest participation index.

30.   The method of claim 26, wherein determining a hamming distance for node pairs within the network includes XORing the node pair members.

IBM Docket No. SJO920010172